DOCUMENT RESUME

ED 043 020 AL 002 583

AUTHOR	Fodor, Jerry
TITLE	Psychology and Language.
INSTITUTION	International Society for Rehabilitation of the Disabled, New York, N.Y.
SPONS AGENCY	Peace Corps (Dept. of State), Washington, D.C.
PUB DATE	[30 Jun 70]
NOTE	33p.; From the "Final Report on Peace Corps Language Coordinators Workshop," Rockport, Massachusetts, April 26 to May 16, 1970
EDRS PRICE	EDRS Price MF-$0.25 HC-$1.75
DESCRIPTORS	Animal Behavior, *Behavior Theories, Child Language, Conditioned Stimulus, *Psycholinguistics, *Psychological Studies, *Second Language Learning, Stimulus Behavior, Stimulus Generalization, *Transformation Theory (Language), Verbal Stimuli

ABSTRACT

The greater part of this paper is dedicated to a non-technical discussion and criticism of the principles of Skinnerian behaviorism. Various aspects of the theory are examined, and its inability to deal with verbal behavior as a productive and creative activity is asserted. The author's point of view is that expressed by Noam Chomsky in his criticism of Skinner's "Verbal Behavior." Evidence is produced to demonstrate that, contrary to the assertions of behaviorists, innate skills are involved in first language learning. The author stresses that although it can therefore be supposed that every child learns his first language in the same way, independently of IQ or culture or training, there is no evidence that second language learning is one particular kind of task approached in the same manner by all learners. For this reason the author believes that there is no basis for believing that there is the possibility of finding a theory that can provide a universally valid technique for language teaching; he stresses instead the pragmatic aspects of language instruction. Finally, in considering how linguistics and psychology can help the language teacher, the author concludes that "Linguistics and psychology have nothing to say about the methods used in foreign language teaching." (FWB)

ED043020

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
OFFICE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL OFFICE OF EDUCATION POSITION OR POLICY.

PSYCHOLOGY AND LANGUAGE

Jerry Fodor

AL 002 583

I will begin by saying something about the Chomsky-Skinner business, and then I'll talk open-endedly about recent results in psycholinguistics proper. I think in a way that the Chomsky-Skinner discussion and the whole operant model of language is a dead issue, and that it is relevant to almost nothing, but people who make materials for teaching you to teach language and people who make gadgets --controlled environments, as they are sometimes called-- for language teaching all claim that they're based on scientific principles. The scientific principles they have in mind are usually the ones I'll be talking about. As I say, I think it all comes to very little, but it's worthwhile knowing the jargon.

Skinnerian psychology is based on a sort of modular system for running psychological experiments. The modular system consists, first of all, of something called a Skinner box. In the box you put a rat, or a pigeon later on in the development of this system. The box is closed, so that the experimenter is not allowed to watch the organism behave, since if he were

allowed to watch the organism behave, that might bias his views on what the behavior was like. So the rat's behavior is recorded electronically.

You begin with a rat in this box and nothing else --just a rat in a completely neutral environment. You start out by making the rat hungry, by subjecting him to about 24 hours of food deprivation. Then you run into this box a bar; under the bar is a trap that food pellets can drop into. Outside the box there is a dispenser for food pellets. And then this system is connected to a graph which will allow you to measure the frequency of the rat's responses as a function of time. Now, in this situation, if the rat is hungry enough, in walking around the box he will at some point happen to bang into the bar. The food dispenser is set up so that whenever he bangs into the bar, a food pellet will drop into the trap. The graph will now allow you to measure the frequency with which he bangs into the bar over time; one coordinate of the graph records the responses and the other coordinate records the time elapsed. What you find is that the rat moves along at zero response level until he hits the bar; at this point the curve goes up a little and maybe flattens out. For a while, little happens, and then the rat hits the bar again, gets another reinforcement, and the curve goes up again, and then levels out. This happens a few times, and then all of a sudden, the rat, as you might say, gets

the idea, and the response wave takes off and flattens out at some fantastically high number. What has happened is that the rat has learned in some sense that he can get a pellet by pressing the bar. And that piece of learning, as it were, shows up objectively in the increase of the response rate, as a function of time.

This is all very straightforward. In fact, if it doesn't happen --if the rat doesn't discover the bar-- then, when nobody's looking, you can open up the box and push the bar down for the rat, and he'll notice that the food pellet comes and will start off that way.

What makes it possible to turn this finding into psychology is the possibility of manipulating various parameters of this situation and looking at how the curve goes as a function of these manipulations. Millions and millions of Ph.D. theses have been done with this sort of modular piece of equipment in exactly that way. For example, you can reinforce the rat, say, every thirteenth time he presses the bar instead of every time he presses the bar, or you can reinforce him only during selected temporal slices (that is, for 5 minutes he get no reinforcement when he presses the bar and then he gets reinforcement any time during thirty seconds, and then for another five minutes he gets no reinforcement, and so on). What we find is that the shape of

the curve is a function of such kinds of variables: the frequency of reinforcement, the hunger of the rat, the amount of reinforcement, the amount of force that is required for him to depress the bar... And since there are in principle indefinitely many values of those variables, you can put out a lot of Ph.D. theses from this kind of design, and that's more or less what has happened.

What occurs in this situation when the rat learns to press the bar under these various conditions is called operant learning, and the whole situation is called an operant paradigm.

It is possible to introduce another wrinkle into this system. If there is a second stimulus, this will provide a further condition which the rat will have to satisfy in order to get reinforced. For example, suppose we put on the wall of the box two lights, a red light and a green light --I don't know whether actually rats are color-blind, but let's assume they're not-- and we set things up so that the rat will be reinforced if he presses the bar when the red light is on, but not when the green light is on. So now reinforcement is contingent on two things --the bar-press and the presence of the red light. In that case, the red light is called the discriminated stimulus: that is, the plus condition, the one for which he gets reinforced. The other one is called the negative discriminated stimulus.

Now, the interesting thing about this --not wildly surprising, but interesting-- is that after a certain amount of training of this kind, a certain correlation develops in the behavior of the rat. If he's reinforced only in the plus condition and not in the minus condition, then what you find is that the blood pressure and behavior of the rat becomes correlated with the presence of the discriminated stimulus; that is, he in fact presses the bar when the light is red, and he doesn't press it when the light is green. This is known as bringing the behavior of the organism under the control of the discriminative stimulus.

There is one further complication that you can introduce into this sort of situation. Suppose that you use a light of x angstrom units for a condition, and you train the rat to respond to the presence of light of this particular shade. Now, in a later trial, after training has taken place, you can test to see what happens if for this red light of x angstroms you substitute a red light of y angstroms, with x not equal to y. What happens is pretty much what you'd expect. Call the light of y angstroms a generalized stimulus, as distinct from the stimulus on which the rat was trained. If you make a graph of the response strength as a function of the quantifiable difference between the generalized stimulus and the discriminative stimulus, a regular behavioral law sets in. If we take the origin to be the discriminative stimulus and we put generalized stimuli at

various distances from the origin, and if this intercept is the trained response rate for the discriminated stimulus, then, roughly speaking, the curve falls off as a function of the distance between the generalized stimulus and the discriminated stimulus. So you can set up a law of behavior which relates the similarity between the originally trained stimulus and the generalized stimulus to the rate of reinforcement, and that states that the extent of generalization, (from the discriminated to the generalized stimulus) varies inversely with the similarity of the stimuli. The more dissimilar they are, the less transfer training and the less generalization you get. This is, of course, a fairly obvious principle. That is, if I teach my dog to come when I whistle in one way, then the probability that he'll come when I whistle in another way is a function of the similarity between the two whistles; the more dissimilar they are, the lower the probability that the dog will come.

Well, that just about exhausts the conceptual content of the Skinnerian system. What I want to do now is to say something about how language looks if you try to reconstruct it on this kind of model.

To apply this model, you have to assume that language consists essentially of a set of habits, and that learning a language consists essentially of learning a set of habits,

because what the organism is doing in the Skinnerian view is learning to respond in a certain way under certain conditions to develop a habit. So the Skinnerians start off by saying that there are two kinds of habits, such that the possession of those two habits constitutes the knowledge of the language. The first is the habit of making a characteristic kind of verbal response in the presence of a characteristic kind of verbal or non-verbal stimulus. That is, you have a discriminated stimulus which may be either verbal or non-verbal, and you have a habit of making a certain kind of verbal response in the presence of that stimulus.

So, for example, I presumably have the habit of saying "crayon" in the presence of a crayon, and the habit of saying "glass" in the presence of a glass. As an example of a verbal stimulus, I presumably have the habit of saying "you're welcome" when som ly says "thank you." So, on this analysis, there's a discriminated stimulus --for example, somebody saying "thank you," or, for that matter, the presence of a glass-- and I've learned a certain verbal response to the occurrence of these stimuli: in the first case I say "you're welcome," and in the second I say "glass." So, this is one kind of habit --the habit of making a characteristic verbal response in the presence of a certain kind of verbal or non-verbal stimulus.

The second kind of habit that is assumed to go into knowing a language is that of making a characteristic non-verbal response in the presence of a certain kind of verbal stimulus. For example, somebody says, "Do you have a light?" and I have the habit of taking out my lighter and lighting it. Or someone says, "Pass the salt," and I have the habit of passing the salt if there happens to be any around. (Or, for that matter, presumably if someone says, "Where were you born?" I have the habit of saying "New York.") So that's a collection of habits which consists of a characteristic disposition to provide verbal responses to certain non-verbal stimuli, and non-verbal responses to certain verbal stimuli. And the suggestion is that when I have developed this set of habits, I know English. That's what it is to me to know English; namely, to have the appropriate habits at these times.

This view of language can be mapped onto the conceptual apparatus of the Skinnerian system. Consider first the habit of producing a certain verbal response to a certain non-verbal stimulus. What the Skinnerians' story has to say as a first approximation is that I got that habit by being reinforced when I produced the appropriate non-verbal response when the non-verbal stimulus was in the area. So, when I was a little boy, I happened to say "glass" when a glass was around, and my Mommy patted me on the head, and that's how I picked up the habit of

saying "glass" when a glass is around. Now of course that can't be exactly true, nor does Skinner suppose it is, because the probability of my happening to make the verbal gesture "glass" in the environment of a glass is presumably negligibly small, when I'm a 1½-year-old baby. Or if my grandfather comes into the room, I have to say, "That's good old grandfather," if I'm going to get reinforced. But the probability of an untrained 1½-year-old baby happening to make the noise, "That's good old grandfather," is presumably infinitely small.

So it can't be that one merely sits around and waits for the baby to say, "That's good old grandfather," in the way that one sits around and waits for the rat to press the bar. This requires a further notion: operant shaping (this is how you get a pigeon to play tennis). Suppose that instead of waiting for the rat in the Skinner box to press the bar, we put him in a cage and he wanders around in the general direction of the bar; what we want him to do is to press the bar. So what we do is everytime he heads in the direction of the bar, we give him a pellet. We rapidly find that the frequency with which he heads toward the bar will increase. Now we change the condition of reinforcement, so that he only gets a pellet when he gets within a certain area fairly close to the bar; and we find that the frequency with which he gets within that area of the bar increases. Now we change the schedule of reinforcement again, so that he gets reinforced only when he's within a certain area

even closer to the bar. By gradual shaping of his behavior, we eventually get him to the position where we want him --he bangs the bar; and we shift the schedule of reinforcement so that he is reinforced only when he actually hits the bar. So now we can have a piece of behavior that is not in fact originally exhibited at any high frequency at all, becoming quite regular.

Suppose we wanted to get the rat to stand on its head --I doubt if he can, but suppose we wanted to. Well, presumably he <u>never</u> stands on his head in his normal ecology, because that's just not one of his tricks. So what we do is give him a pellet first when he puts his head down a little; then, when we've got him putting his head down very frequently, we give him a pellet only when he happens to hit the floor when he puts his head down. And so forth and so on, until by some set of successive approximations, we give him a pellet only when he stands on his head, in case he ever happens to do so.

Well, it's exactly this kind of operant shaping that goes on when Mommy gets the baby to say "glass." When the baby is brought into an environment in which there happens to be a glass, in the course of his random babbling he will sometimes make sounds like "NGMPPP" or"GIHH." These sounds would be set up to be reinforced in the presence of a glass. That is, Mommy will pat him on the head and give him an M&M or whatever mothers do,

when he makes that sound. When he is regularly making this sound in the presence of a glass, we can shift the criteria of reinforcement, so that we only reinforce him for those "GIHH" sounds which are very close to the first few sounds of "glass." If we keep shifting our criteria of reinforcement to response that are successively closer to the correct pronunciation of "glass," over a long enough period of time, presumably we will shape the baby's behavior, so that he now says "glass" and nothing else in the presence of the discriminated stimulus of a glass. (You can play this another way, and some Skinnerians are inclined to. You can claim that on top of the apparatus described above, the child also has an innate tendency to imitate. In this case, one could assume that the parent says "glass" when a glass is around and the child tries to imitate that, and then the quality of his imitation and his tendency to say "glass" are shaped by reinforcement. This gets you off the ground much faster.)

The other set of habits --making a specific kind of non-verbal response to a specific kind of verbal stimulus-- are treated in a similar fashion. For example, when I was three years old, somebody said, "Pass the cigarettes," and I passed the cigarettes and got patted on the head, and this happened with sufficient frequency that I developed the habit of producing the non-verbal response of passing the cigarettes in the presence

of the verbal discriminated stimulus of somebody saying "Pass the cigarettes."

Now here is where generalization comes in. You might suppose, on the basis of the kind of account presented above, that I would learn to say "glass" in the presence of a glass, but not when this same glass is turned upside down. But that is covered by the generalization clause, which says that, given that I've been trained to say "glass" in the presence of a glass, the probability of my saying "glass" in the presence of an upside-down glass is a function of the similarity between the two, in some rather grossly undefined sense of similarity.

The principle of generalization also accounts for how one produces a response to stimuli which are distinct but similar. For example, I'll pass the cigarettes whether I'm asked to do so by an adult, who produces a certain kind of acoustic pattern when he says, "Pass the cigarettes," or by a three-year-old, who produces a different kind of acoustic pattern. This is accounted for by the principle of generalization, which says that my training transfers from one stimulus to other stimuli of characteristically similar kinds.

Well, that is essentially --with certain mildly offensive aberrations, perhaps-- what the Skinnerians' account of language

looks like. What you have is a conceptual mechanism of saying how language is put together, for saying how the set of habits characteristic of having a language are assimilated --given that you assume that having that set of habits is what is involved in having language. And, in a sense, there is just nothing else to say. You've explained what understanding a sentence is --that is, being disposed to produce the appropriate non-verbal response, given the sentence that has been uttered; and you've explained what speaking the language is --that is, being disposed to produce certain verbal operants, given the appropriate input stimulation. So, now that you've said what it is to understand the language and to speak it, you've said everything there is to say.

(Suppose you wanted to apply this kind of system to a language learning situation. You would put somebody in the following condition: you show a glass on a screen or something, and through a loudspeaker you say "glass." If he says "glass" properly, you give him an M&M or something; or you wait around for him to say something that sounds like "glass" and then give him an M&M, and eventually he ends up calling a glass a glass. Now, in a certain sense, I can't really see what else you can do. If you want to teach somebody what a glass is called, and the game is so played that you're not allowed to say, "Look, this is called 'glass,'" for some methodological reason --because it's

not considered scientific or something-- then about all you can do is wait around for him to say "glass" and tell him "you're right" when he says it. This either works or it doesn't.)

But what I want to suggest is that the view of language on which this kind of story is based is just hopeless. It seems immediately obvious that these kinds of habits aren't really what's at issue. Let me just make a couple of crude points and then go into this in detail. It is perfectly obvious that in fact I don't have the habit of saying "glass" whenever a glass is presented, because I don't say "glass" whenever a glass is presented. There are billions of things of all kinds in the environment --people, faces, one's nose, and so forth-- which are allegedly stimulus inputs but which you don't name. Now, that's a sort of banal point, but it's very important. When you talk a language, your behavior vis-a-vis the stimuli in your environment is very different from the rat's behavior vis-a-vis the bar-press. Once the rat has learned its trick, it presses the bar whenever there's a chance that this will lead to a reward. Language doesn't work like that. There simply isn't any specifiable set of response habits such that one can rely upon a speaker producing a certain utterance given that a certain stimulus situation has been presented, in the way that you can rely on a rat pressing the bar after a certain period of training.

Now, this is so self-evident that Skinner even noticed it in Verbal Behavior, and what he says is that we can bring verbal behavior under direct stimulus control, in the following way: we take a guy and we bring him into an absolutely bare room, in which the only noticeable object is a pencil 16 feet tall. Now, Skinner says that it's under those circumstances that it's highly probable that the guy will say "pencil." This is probably correct. The question is whether or not one wants to treat that as a paradigm of an instance in which one uses the language. As another paradigm of such an instance, suppose I'm a lot bigger than you are, and I say "Say 'uncle' or I'll break your arm." The probability under those circumstances is very high that if you believe me, you'll say "uncle." Or, I approach you with a large pistol and say, "Say 'uncle' or I'll blow your brains out." And, depending on your view of what my character is, you're very likely to say "uncle" under those circumstances.

But what one notices instantly is that these sorts of circumstances, in which there is a direct correlation between the character of the stimulus situation and the character of the response, are highly atypical of language use. You have to put a person literally under physical duress in order to get such a response, and then the question whether he responds appropriately or not is a function of how much duress he's going to stand for, and not a question of his verbal habits.

In a way, that same thing happens with the rat in the Skinner box; the only way to get the rat to respond reliably is by making him very hungry.

Now, I think that there's a general consideration which underlies this that makes the whole Skinnerian business extremely uninteresting. That is, in any kind of ongoing behavior by any organism --and especially language behavior by people-- the actual output is a function of at least three things: the organism's knowledge of the response system (the rat's knowledge of how to press the bar, a person's knowledge of how to speak the language), plus the organism's utilities, plus what you could call the organism's local beliefs. By the rat's knowledge, in this case, I mean that he has the actual muscular coordinations available to press the bar. He has certain utilities which come out in the characteristic Skinner situations, so that when he's hungry, he behaves reliably in accordance with what we know to be his knowledge. And he has certain local beliefs --that is, what he has learned about the relation between his trained response and getting reinforced. Now, in the Skinnerian situation, you have known values for the rat's knowledge, utilities, and local beliefs. The only belief that the organism could possibly develop in this situation is that pressing the bar leads to reinforcement, because the bar is the only thing in the environment: that is, the only

available induction for him to make is that it's pressing the bar and not something else that leads to reinforcement. Given that these three are known values, you can make extremely precise predictions about behavior.

But now look at the situation vis-a-vis language. If you think of language as literally a correlation between a stimulus situation and a response situation --if, for example, you regard language as a piece of knowledge that you should produce a certain response given a certain input, then you have to look for correlations of this kind, and you don't find them. The reason that you can't find them is intuitively evident--namely, what response a particular person gives to a particular stimulus at a particular time will depend on a variety of things other than his knowledge of the response system, that is, other than his knowledge of the language. It will depend on what his utilities are, what he believes, and so forth, and no psychology has anything to say about these kinds of conditions. So that if I walk up to a guy and say, "What's the time?" and if I can make the usual assumptions about his utilities and his local beliefs --his utilities will include things like politeness, and his local beliefs will include the fact that his watch isn't broken, and so on-- then I can predict what he'll do: he'll look at his watch and say "3 o'clock," or whatever time it is. But if he happens to be drunk or angry, then he's likely to punch

me in the nose. Certainly, whether or not he punches me in the nose has absolutely nothing to do with whether he understands the question "What's the time?" And that's the important part. What his response is can throw exactly no light on the question of what he knows about the language. And that is a way of saying that there is no possibility of construing his knowledge of the language in terms of a disposition to produce a characteristic response in a characteristic stimulus situation.

This is the critical mistake: to suppose that what a speaker knows is that he is supposed to give a certain response in the appropriate stimulus situation. What I think is really going on is that the person's knowledge of the language has some wildly different form from that which linguists, psycholinguists, and philosphers try to study; and the response is a function of what I think are actually a lot of different variables. That the response is in English rather than Hindustani is obviously a function of what he knows about the language, but which response he gives in a particular instance will depend not only on his knowledge of the language, but also on many other factors, including his utilities and beliefs.

If you read Chomsky's review of Verbal Behavior, and I highly recommend that you do, you get the feeling that there is a radical difference between language, which has the properties

of productivity and stimulus freedom, and what the rat does in a Skinner box. What I'd be inclined to say is that Skinner even gives a bad analysis of what the rat does. What we've done is put the organism in a situation where the only strategies available are brute force strategies. Any organism behaves differently when you watch it behaving in its natural ecology, from the way it behaves when you watch it in a Skinner box. This is clearly an impoverished situation. The only intelligence that the organism may have is reduced by the brute exigencies of the situation to an induction about the results of pressing a bar. If, on the other hand, you watch a rat building a nest, which they do, you get a quite different picture of the capacities of the animal.

I think, in fact, that in this respect there's no difference between a rat and a person, except that in the Skinnerian situation the rat is extremely hungry and is given a very limited kind of situation in which to form his local beliefs. That is, if you put a rat into this completely controlled kind of environment, in which a) he wants food, and b) there's only one kind of manipulation he can perform on the environment --he can press a bar-- then the only conceivable induction he can make is that he gets a pellet by pressing the bar. It may be that there are specific differences in the kinds of strategies that various organisms will try for this situation, but I'm

inclined to doubt it. A super-smart rat might think to himself, "Well, there's only one thing I can do in this situation, so I'll go do it," and a characteristically dumb rat might bang around until he runs into the bar.

Looking at the rat's response curve for this situation, you'd expect that if he were doing it by trial and error, you would get a smooth, slow approximation. In fact, that's not what happens at all. His response doesn't increase at a smooth rate. What happens is that he does nothing at all for a while, and then he makes a couple of responses, and then the curve really takes off. The suggestion is that the rat is doing exactly what you're doing --he sees what's going on. He sees that there's some connection between pressing the bar and getting food, and he learns to handle the situation. So I think that it's dubious that there's any difference between the way a person would address this situation and the way a rat does.

What I wish to argue now is that the basic reason that the Skinnerian system fails is that the principle of generalization has no content. Just as you have stimulus generalization, in which, for example, the habit of saying "glass" to a right-side-up glass extends to an upside-down glass, so also you have response generalization. For instance, if you train a rat in

a box with the bar in one place and then put him in another box, with the bar in a different place, it will take him much less time to learn the task. Now, it is certainly true that for any of the skills -- language or what the rat does in the Skinner box -- the productivity of the skill is extremely important. Here there is a real fact to be explained, namely, that the organism, having learned one thing, transfers his training to something quite different. And the principle of generalization doesn't explain it.

There is a well know experiment of a T maze, in which there's a goal box in one direction which is plus, and a goal box in the other direction which is minus, and what you do is turn the rat around like this, and after a while, he does that. The rat starts out here, and he explores for a while, and eventually he finds the food here.

Now what you do is flood the maze, so that the organism can't run it, but must instead swim in order to reach the goal. This entails a completely new set of motor responses. What happens is that the rat, previously trained to run the unflooded maze, successfully swims the flooded maze. Now, one might say that this is an instance of response generalization. But

notice what question is being begged. You have to say that swimming to the right and running to the right are somehow more similar to each other than running right and running left, because, after all, this response generalizes to swimming to the right, but not to running to the left. What's the principle of similarity that allows you to do that? How can you say that running right and swimming right are more similar than running right and running left, when the latter pair in fact uses exactly the same motor system, whereas the former pair uses different motor systems?

Exactly the same kind of question arises about language. You might want to say that language is productive -- that you can understand sentences you've never heard before. Suppose you've been trained in the sentence

(1) Mary went to school.

How do you know that another sentence, which may be quite different, like

(2) The boy who I used to know when I went golfing in Chicago got sick

is also a sentence, given that you weren't specifically trained in it? Well, you might say that you just generalize from the former to the latter. But what is the principle of generalization that says that (2) is relevantly similar to (1) and so the training transfers, but the sentence

(3) School to went Mary

is not relevantly similar? If you don't spell out the notion of similarity, the principle of generalization is vacuous. On the other hand, if you do try to spell out the principles that underlie the fact that (2) but not (3) is similar to (1) --namely, that (1) and (2) are grammatical sentences and that (3) isn't-- then you get a very complex set of grammatical rules; what you get in effect is a generative grammar. Thus, in a sense, the principle of generalization doesn't buy anything. There are a vast set of structures which the organism employs in determining what is similar. Using the principle of generalization to answer one question --how you transfer training from the structures you've been trained on to new objects-- just opens up another question: how does the organism learn those principles which determine which responses are similar and which are not?

Similarly, suppose that I've been trained to call a non-filter cigarette "cigarette." What is the principle of similarity that allows me to generalize this response, so that

I will also call a filter cigarette "cigarette," but I will not call a piece of chalk "cigarette"? Certainly, a piece of chalk is similar in some respects to a cigarette; the principle of generalization has nothing to say about why it is not relevantly similar.

Or, suppose you train an organism to respond positively to a certain note being played on the piano. If you look at the way he generalizes the responses, you may guess that he will respond positively to similar notes and not to dissimilar notes. What determines whether a note is similar to the note to which the organism has been trained to respond? Well, if you define similarity in terms of simple frequency, so that notes of similar pitch will be relevantly similar, you turn out to be wrong. If you define similarity in terms of an octave, so that notes which have different frequency but are in the same key are considered relevantly similar, you turn out to be right. But the principle of generalization, since it is not equipped to tell you what is similar and what isn't, has told you nothing.

Or, suppose you train an organism to press a bar when you show him a red triangle, but not when you show a green circle. What will happen now if you show a green triangle or a red circle? It is clear that the principle of generalization is of no help here; it can make no prediction, since you don't

know what similarity means to the organism in question. (By the way, there is no general answer in this case; whether the organism will respond to the green triangle or red circle depends on many factors: its species, maturity, prior training, etc.)

If you train a person to respond positively to the word "toy," which of the following words will he tend to respond positively to: "toy" or "man"? Of course, the principle of generalization cannot predict, since it provides no function to tell you how the person arranges stimuli into similarity classes (the answer is that children in this case tend to respond positively to "toy," adults to "man").

Chomsky, in his review of Verbal Behavior, says that language has two properties that cannot be accounted for by Skinner. The first is productivity. If language is a habit, then it is a habit with infinitely many possible expressions. Since you can't train an organism to do infinitely many tasks, at some point training has to stop and the organism has to take over and generalize. There is no account for this in the Skinnerian model.

The second property of language that is left unaccounted for in the Skinnerian model is stimulus freedom. That is, what

one says is not a simple function of one's input. This point has been largely misunderstood. Chomsky has been taken as asserting that language behavior is not deterministic--that it is not physically determined. This assertion may or may not be true, depending on your point of view, but it was not the point of Chomsky's remark. What he intended was that the response is a function not only of the input but also of the internal states of the organism. This, of course, is true not just of language, but of just about anything.

I now wish to sketch some lines of evidence that first language learning in humans involves innate skills. To do this, I want to turn to an aspect of the language problem that was hardly noticed by the Skinnerians. When you hear a word, part of what you do in understanding it is figure out which word you're hearing. You are able to do this across many differences among speakers --speed of speech, whether a male or a female is talking, and so forth. How, in spite of these differences, do you manage to identify a given word as that particular word, or a given sentence as that particular sentence? What the Skinnerians said is that you are trained to recognize certain critical acoustic features associated with a sentence that are in a one-to-one correspondence with the linguistic analysis. In effect, the Skinnerians said that speech perception is no problem, because speech perception is equivalent to

responding to acoustic cues.

However, it is very unlikely that this is true. After 25 years, we still haven't been able to build a machine which can process the wave forms containing acoustic cues and interpret them as phonetic symbols. If speech recognition were just a matter of responding to acoustic cues, we should be able to build such a machine.

The reason that we haven't been able to build a machine to process continuous speech is the following: the acoustic representation of a particular consonant is determined almost entirely by the vowel that it is attached to. Now, it is possible in some cases to have two distinct consonants, one of them before one vowel, and the other before a different vowel, where these distinct consonants have the same acoustic representation. This situation does not lead to confusion in your perception of these sounds; they are perceptually distinguishable, but they are indistinguishable acoustically if you look at their representations on a sound spectrogram.

The reason that the two consonant sounds are perceptually distinguishable is that in your head you have for each sound a phonetic matrix of distinctive features. For example, you can describe the sound u in terms of various features, such

as roundedness, voicing, consonantality, continuousness, and so forth, where each of these features has a particular value. Each distinctive feature matrix is in effect a set of discrete instructions at the phonetic level. Because of the limitations of your vocal apparatus, these discrete instructions get smeared all over the place when it comes to pronouncing this sound in a word. Thus what you do in pronouncing this sound is go from a digital representation (the discrete phonetic matrix in your head) to an analog representation (that is, a continuous system employing your vocal apparatus, which is inertial).

Actually, the situation is somewhat worse than this. If you pronounce the words "sloop" and "slip" you will notice that some rounding occurs at the beginning of the word "sloop" even though the significant rounding occurs later, in the vowel. So what you have is a situation where the actual output --for example, the rounding of the s in sloop-- is determined not by the discrete set of instructions for the sound s, but rather by a feature of a subsequent sound, in this case, the vowel oo. So the instructions are not even transmitted sequentially in producing the sounds.

Thus, to understand speech, you have to do two things: first, you have to compensate for the fact that the vocal

apparatus is a noisy analog transducer, and second, you need non-sequential strategies for processing the signal. This is why we haven't been able to build a machine to do the job.

How, then, does the child learn to perform the seemingly complex operations involved in processing speech? The above considerations make it difficult to maintain that he is <u>trained</u> to do so. So the answer we ought to consider is that the infant has innately available to him a theory of relating discrete phonetic matrices to a continuous inertial system, and a theory of non-sequential processing strategies. In fact, there is growing evidence that the child's perceptual analysis, at least by the age of 3 or 4 months and perhaps even earlier, is essentially the same as the adult's. The idea is, then, that children are specially built to do this, rather than having to learn these complex theories, which at the present state of our knowledge we can't even explain.

So here we are dealing with a very highly specialized organism, which starts out with a great deal of highly structured information on what kinds of strategies to use on its input. This view has two immediate consequences for language teaching. First, teaching a child how to speak his first language is about as useful as teaching a bird to fly. In fact, there is even some amusing evidence that if you try to train the child, you'll

just slow him up. That depends on <u>how</u> you train him, but the things that have been tried so far have had negative effects. It appears that exposure to people talking language under normal circumstances is just optimal for learning a first language. Second, if these innate abilities are no longer available to the child after some age, say when he reaches adolescence --and there is some evidence for this-- then you can't draw any inferences for second language learning from first language learning. If this is true, then the only implication you can draw from this theory to language teaching is that this theory has nothing to say about language teaching. I think that this may be exactly right. All that you read in the journals about scientific techniques for second language teaching is absolute nonsense. There is nothing in current linguistics or psychology that gives any indication about second language teaching.

In fact, the situation might be worse than that. There is real reason to suppose that every child learns his first language in the same way, independently of IQ, or cultures, or training, or what language it is, or anything of the sort. But there is no reason to suppose that second language teaching is one particular kind of task. Now, if it is true that John learns a second language one way, and Mary another way, and Peter does it 14 different ways, then the answer to the question "How do

you teach a language?" might be "Any way you like." Now, I'm not necessarily saying that this is necessarily true; what I <u>am</u> saying is that, on the basis of the current work in psychology, there is no reason to believe that it's not true.

There is a belief that seems quite prevalent among language teachers, that they need to have a theory in order to do their job. But I don't agree. For example, if you're going to teach the skill of driving, you're not going to go about it by having a general theory of internal combustion engines or something and then deducing theorems about driving from it. What you'll do is buy some automobiles and take your students out on the road and say, "Drive.--I'll stop you if you do something wrong." If someone claims to have a fool-proof, scientifically based theory of driving, he's a liar. But the lack of a theory doesn't prevent driving instruction from going on. Similarly, in my remarks here, I'm not suggesting that you stop teaching languages because there is no theory of language teaching. What I'm suggesting is that we learn to live with the fact that there is no theory of language teaching at present. In the absence of a theory, what you have to do is use your head. You make it up as you go along, and you try to do the best you can.

Someone might suggest that the reason that different teachers have different degrees of success is attributable to

the sort of theory within which they are operating, but it may well be the case that the theories are independent of the results that follow. There is certainly no known correlation between the theory and the effectiveness of its application.

I wish to stress the fact that just because linguistics and psychology have nothing to say about the methods used in foreign language teaching, it does not follow that there may not be some methods of teaching that are better than others. I don't think that these methods will come out of the work being done in linguistics or psychology, but this is not to say that such methods can't be found.

Let me conclude with a general remark. Classically, in the view of people like Plato and Spinoza, education pretty much consisted in applied ethics. That is, the question you asked was, "What should kids learn and what shouldn't they learn, given that they can't learn everything?" Fairly recently, since the time, say, of Dewey, a new view of education came into prominence--namely, that education is more or less a branch of applied psychology, that it should be concerned primarily with technique. I don't agree with this view at all. A lot of classroom procedures have been proposed for which there is not the least bit of evidence that they are preferable to other conceivable procedures. For example, there is an

unquestioned belief in the importance of the student-teacher ratio, but not a shred of evidence has been offered to show that this has anything to do with what the students learn. In fact, it is a completely open question whether any teacher is better than none. Basically I believe that there is no reason to think that there is a theory that can provide a technique for teaching, and that an undue amount of time has gone into searching for such a theory.